

Patent Application

for

METHOD OF UTILIZING A SINGLE UNIFORM
RESOURCE LOCATOR FOR RESOURCES
WITH MULTIPLE FORMATS

by

Nils B. Lahr

This application claims the benefit of U.S. provisional application Serial No. 60/178,751, filed January 28, 2000.

Cross Reference to Related Applications:

Related subject matter is disclosed in co-pending U.S. patent application of Nils B. Lahr et al., filed September 28, 1998, entitled "Streaming Media Transparency" (attorney's file IBC-P001); in co-pending U.S. patent application of Nils B. Lahr, filed
5 even date herewith, entitled "Method and Apparatus for Encoder-Based Distribution of Live Video and Other Streaming Content" (attorney's file 39512A); in co-pending U.S. patent application of Nils B. Lahr, filed even date herewith, entitled "A System and Method for Rewriting Media Resource Request and/or Response Between Origin Server and Client" (attorney's file 39511A); in co-pending U.S. patent application of
10 Nils B. Lahr, filed even date herewith, entitled "Method and Apparatus for Client-Side Authentication and Stream Selection in a Content Distribution System" (attorney's file 39505A); in co-pending U.S. patent application of Nils B. Lahr, filed even date herewith, entitled "Method and System for Distributed Data Mining and Analysis for Networks" (attorney's file 39510A); in co-pending U.S. patent application of Nils B.
15 Lahr et al., filed even date herewith, entitled "A System and Method for Mirroring and Caching Compressed Data in a Content Distribution System" (attorney's file 39565A); in co-pending U.S. patent application of Nils B. Lahr, filed even date herewith, entitled "A System and Method for Determining Optimal Server in a Distributed Network for

Serving Content Streams" (attorney's file 39551A); and in co-pending U.S. patent application of Nils B. Lahr, filed even date herewith, entitled "A System and Method for Performing Broadcast-Enabled Disk Drive Replication in a Distributed Data Delivery Network" (attorney's file 39564A); the entire contents of each of these applications being expressly incorporated herein by reference.

Field of the Invention:

The present invention relates to method of providing and utilizing a single Uniform Resource Locator for a resource which may be provided in a plurality of different formats and via different servers via the Internet.

Background of the Invention:

A Uniform Resource Locator (URL) is an address of a resource or file accessible on the Internet. The URL contains the name of the protocol required to access the resource, a domain name that identifies a specific computer or server on the Internet, and a hierarchical description of a file location on that computer or server. Presently, a single URL is associated with a single resource available on the Internet. There are many different types of Internet resources or client-server applications that are provided in different formats. If a resource is hosted in multiple formats, then a separate link to each format must be provided. This is problematic in that it can greatly complicate content management and create extra work. Moreover, as more formats are introduced, the problem increases.

Summary of the Invention:

In accordance with the present invention, a method and system are provided to supply content hosting for resources having multiple formats by providing a single URL for a resource having multiple formats.

In accordance with an aspect of the present invention, a testing component uses information regarding the client and service provider provided in the metafile corresponding to the request to determine the type of client from which the request originated and to redirect the client to a server that can service the request.

In accordance with another aspect of the present invention, the testing component is implemented in a server.

Brief Description of the Drawings:

5 These and other objects, advantages and novel features of the invention will be more readily appreciated from the following detail description when read in conjunction with the accompanying drawing, in which:

 Fig. 1 illustrates a client-side testing and stream selection device constructed in accordance with an embodiment of the present invention;

10 Fig. 2 illustrates an Internet broadcast system for streaming media constructed in accordance with an embodiment of the present invention;

 Fig. 3 is a block diagram of a media serving system constructed in accordance with an embodiment of the present invention;

15 Fig. 4 is a block diagram of a data center constructed in accordance with an embodiment of the present invention;

 Fig. 5 illustrates data flow in a Internet broadcast system for streaming media constructed in accordance with an embodiment of the present invention;

20 Figs. 6, 7 and 8 illustrate acquisition, broadcasting and reception phases, respectively, employed in a Internet broadcast system for streaming media constructed in accordance with an embodiment of the present invention; and

 Fig. 9 illustrates transport data management in a Internet broadcast system for streaming media constructed in accordance with an embodiment of the present invention.

25 Throughout the drawing figures, like reference numerals will be understood to refer to like parts and components.

Detailed Description of the Preferred Embodiments:

30 In accordance with the present invention, a testing component 27 is provided between a client 20 and a server 25, as shown in Fig. 1. The testing component 27 is operable to receive a request from the client 20 for access and to selectively establish a network connection 29 with the server 25. As described below, the request is analyzed

to determine information regarding the client with which to logically redirect the client to a server for serving a particular media type.

The testing component 27 can be implemented as an add-on software and/or hardware component provided that it can be configured to selectively establish or intercept a network connection 29 between the client 20 and the server 25. The client-server relationship is described herein for illustrative purposes in connection with a content distribution system. An overview of an exemplary content distribution system 10 follows. The system 10 employs data transport protocols to implement the testing component 27 in accordance with the present invention. It is to be understood that implementation of the invention is not limited to the architecture of the illustrative system 10 described herein.

1. System Component Overview

With reference to Fig. 2, a system 10 is provided which captures media (e.g., using a private network), and broadcasts the media (e.g., by satellite) to servers located at the edge of the Internet, that is, where users 20 connect to the Internet such as at a local Internet service provider or ISP. The system 10 bypasses the congestion and expense associated with the Internet backbone to deliver high-fidelity streams at low cost to servers located as close to end users 20 as possible.

To maximize performance, scalability and availability, the system 10 deploys the servers in a tiered hierarchy distribution network indicated generally at 12 that can be built from different numbers and combinations of network building components comprising media serving systems 14, regional data centers 16 and master data centers 18. The system also comprises an acquisition network 22 that is preferably a dedicated network for obtaining media or content for distribution from different sources. The acquisition network 22 can operate as a network operations center (NOC) which manages the content to be distributed, as well as the resources for distributing it. For example, content is preferably dynamically distributed across the system network 12 in response to changing traffic patterns in accordance with the present invention. While only one master data center 18 is illustrated, it is to be understood that the system can

employ multiple master data centers, or none at all and simply use regional data centers 16 and media serving systems 14, or only media serving systems 14.

An illustrative acquisition network 22 comprises content sources 24 such as content received from audio and/or video equipment employed at a stadium for a live
5 broadcast via satellite 26. The broadcast signal is provided to an encoding facility 28. Live or simulated live broadcasts can also be rendered via stadium or studio cameras, for example, and transmitted via a terrestrial network such as a T1, T3 or ISDN or other type of a dedicated network 30 that employs asynchronous transfer mode (ATM) or other technology. In addition to live analog or digital signals, the content can include
10 analog tape recordings, and digitally stored information (e.g., media-on-demand or MOD), among other types of content. Further, in addition to a dedicated link 30 or a satellite link 26, the content harvested by the acquisition network 22 can be received via the Internet, other wireless communication links besides a satellite link, or even via shipment of storage media containing the content, among other methods. The encoding
15 facility 28 converts raw content such as digital video into Internet-ready data in different formats such as the Microsoft Windows Media (MWM), RealNetworks G2, or Apple QuickTime (QT) formats. The system 10 also employs unique encoding methods to maximize fidelity of the audio and video signals that are delivered via multicast by the distribution network 12.

20 With continued reference to Fig. 2, the encoding facility 28 provides encoded data to the hierarchical distribution network 12 via a broadcast backbone which is preferably a point-to-multipoint distribution network. While a satellite link indicated generally at 32 is used, the broadcast backbone employed by the system 10 of the present invention is preferably a hybrid fiber-satellite transmission system that also
25 comprises a terrestrial network 33. The satellite link 32 is preferably dedicated and independent of a satellite link 26 employed for acquisition purposes. The tiered network building components 14, 16 and 18 are each equipped with satellite transceivers to allow the system 10 to simultaneously deliver live streams to all server tiers 14, 16 and 18 and rapidly update on-demand content stored at any tier. When a satellite link
30 32 is unavailable or impractical, however, the system 10 broadcasts live and on-demand content though fiber links provided in the hierarchical distribution network 12. Where

the feed is pulled from in case of a failure is based on a set of routing rules that include priorities, weighting, and so on. In other words, the feed is pulled in a manner similar to the way routers currently operate, but at the actual stream level.

5 The system 10 employs a director agent to monitor the status of all of the tiers of the distribution network 12 and redirect users 20 to the optimal server, depending on the requested content. The director agent can originate, for example, from the NOC/encoding facility 28. The system employs an Internet Protocol or IP address map to determine where a user 20 is located and then identifies which of the tiered servers 14, 16 and 18 can deliver the highest quality stream, depending on network
10 performance, content location, central processing unit load for each network component, application status, among other factors. Cookies and data from other databases can also be employed to facilitate this system intelligence.

Media serving systems 14 comprise hardware and software installed in ISP facilities at the edge of the Internet. The media serving systems preferably only serve
15 users 20 in its subnetwork. Thus, the media serving systems 14 are configured to provide the best media transmission quality possible because the end users 20 are local. A media serving system 14 is similar to an ISP caching server, except that the content served from the media serving system is controlled by the content provider that input the content into the system 10. The media serving systems 14 each serve live streams
20 delivered by the satellite link 32, and store popular content such as current and/or geographically-specific news clips. Each media serving system 14 manages its storage space and deletes content that is less frequently accessed by users 20 in its subnetwork. Content that is not stored at the media serving system 14 can be served from regional data centers.

25 With reference to Fig. 3, a media serving system 14 comprises an input 40 from a satellite and/or terrestrial signal transceiver 43. The media serving system 14 can output content to users 20 in its subnetwork or control/feedback signals for transmission to the NOC or another hierarchical component in the system 10 via a wireline or wireless communication network. The media serving system 14 has a central processing unit 42
30 and a local storage device 44. A file transport module 136 and a transport receiver 144, which are described below with reference to Fig. 8, are provided to facilitate reception

of content from the broadcast backbone. The media serving system 14 also preferably comprises one or more of an HTTP/Proxy server 46, a Real server 48, a QT server 50 and a WMS server 52 to provide content to users 20 in a selected format. The media serving system can also support Windows and Real caching servers, allowing direct
5 connections to a local box regardless of whether the content is available. The content in the network 12 is then located and cached locally for playback. This allows for split live feeds by a local media serving system 14 regardless of whether is being sent via a broadcast or feed mechanism. Thus, pull splits from a media serving system 14 are supported, as well as broadcast streams that are essentially push splits with forward
10 caching. Also, the database 44 and file system 136 can be local or remote, depending on where the media serving system 14 is installed.

The regional data centers 16 are located at strategic points around the Internet backbone. With reference to Fig. 4, a regional data center 16 comprises a satellite and/or terrestrial signal transceiver, indicated at 61 and 63, to receive inputs and to
15 output content to users 20 or control/feedback signals for transmission to the NOC or another hierarchical component in the system 10 via wireline or wireless communication network. A regional data center 16 preferably has more hardware than a media serving system 14 such as gigabit routers and load-balancing switches 66 and 68, along with high-capacity servers (e.g., plural media serving systems 14) and a storage
20 device 62. The CPU 60 and host 64 are operable to facilitate storage and delivery of less frequently accessed on-demand content using the servers 14 and switches 66 and 68. The regional data centers 16 also deliver content if a standalone media serving system 14 is not available to a particular user 20. The director agent software preferably continuously monitors the status of the standalone media serving systems 14 and
25 reroutes users 20 to the nearest regional data center 16 if the nearest media serving system 14 fails, reaches its fulfillment capacity or drops packets. Users 20 are typically assigned to the regional data center 14 that corresponds with the Internet backbone provider that serves their ISP, thereby maximizing performance of the second tier of the distribution network 12. The regional data centers 14 also serve any users 20 whose ISP
30 does not have an edge server.

The master data centers 18 are similar to regional data centers 16, except that they are preferably much larger hardware deployments and are preferably located in a few peered data centers and co-location facilities, which provide the master data centers with connections to thousands of ISPs. With reference to Fig. 4, master data centers 18 comprises multiterabyte storage systems (e.g., a larger number of media serving systems 14) to manage large libraries of content created, for example, by major media companies. The director agent automatically routes traffic to the closest master data center 18 if a media serving system 14 or regional data center 16 is unavailable. The master data centers 18 can therefore absorb massive surges in demand without impacting the basic operation and reliability of the network.

Transmissions can occur out of the data centers 16 and 18. In the case of the satellite 32, however, transmissions can also be implemented by taking what is being received and routing a copy thereof directly to the uplink system without first passing through the media serving systems 14.

2. System Operation Overview

With reference to Fig. 5, the Internet broadcast system 10 for streaming media generally comprises three phases, that is, acquisition 100, broadcasting 102 and receiving 104. In the acquisition phase 100, content is provided to the system 10 from different sources such as Internet content providers (ICPs) or event or studio content sources. As stated previously, content can be received from audio and/or video equipment employed at a stadium for a live broadcast. The content can be, for example, live analog signals, live digital signals, analog tape recordings, digitally stored information (e.g., media-on-demand or MOD), among other types of content. The content can be locally encoded or transcoded at the source using, for example, file transport protocol (FTP), MSBD, or real-time transport protocol/ real-time streaming protocol (RTP/RTSP). The content is collected using one or more acquisition modules 106, which are described in more detail below in connection with Fig. 6. The acquisition modules 106 represent different feeds to the system 10 in the acquisition network 12 and can be co-located or distributed. Generally, acquisition modules 106 can perform remote

transcoding or encoding of content using FTP, MSBD, or RTP/RTSP or other protocols prior to transmission to a broadcaster 110 for multicast to edge devices and subsequent rendering to users 20 located relatively near to one of the edge devices. The content is then converted into a broadcast packet in accordance with an aspect of the present invention. This process of packaging packets in a manner to facilitate multicasting, and to provide insight at reception sites as to what the packets are and what media they represent, constitutes a significant advantage of the system 10 over other content delivery systems.

Content obtained via the acquisition phase 100 is preferably provided to one or more broadcasters 110 via a multicast cloud or network(s) 108. The content is unicast or preferably multicast from the different acquisition modules 106 to the broadcasters 110 via the cloud 108. As stated above, the cloud 108 is preferably a point-to-multipoint broadcast backbone. The cloud 108 can be implemented as one or more of a wireless network such as a satellite network or a terrestrial or wireline network such as optical fiber link. The cloud 108 can employ a dedicated ATM link or the Internet backbone, as well as a satellite link, to multicast streaming media. The broadcasters 110 are preferably in tier 120, that is, they are master data centers 18 that receive content from the acquisition modules 106 and, in turn, broadcast the content to other receivers in tiers 116, 118 and 120.

During the broadcasting phase 102, broadcasters 110 operate as gatekeepers, as described below in connection with Fig. 7, to transmit content to a number of receivers in the tiers 116, 118 and 120 via paths in the multicast cloud 108. The broadcasters 110 support peering with other acquisition modules indicated generally at 112. The peering relationship between a broadcaster 110 and an acquisition module 112 is via a direct link and each device agrees to forward the packets of the other device and to otherwise share content directly across this link, as opposed to a standard Internet backbone.

During the reception phase 104, high-fidelity streams that have been transmitted via the broadcasters 110 across the multicast cloud 108 are received by servers 14, 16 and 18 located as close to end users as possible. The system 10 is therefore advantageous in that streams bypass congestion and expense associated with the Internet backbone. As stated previously, the servers are preferably deployed in a tiered hierarchy comprising

media serving systems 14, regional data centers 16 and master data centers 18 that correspond to tiers 116, 118 and 120, respectively. The tiers 116, 118 and 120 provide serving functions (e.g., transcoding from RTP to MMS, RealNet, HTTP, WAP or other protocol), as well as delivery via a local area network (LAN), the Internet, a wireless network or other network to user devices 122 for rendering (e.g., PCs, workstations, set-top boxes such as for cable, WebTV, DTV, and so on, telephony devices, and the like). The tiers in the reception phase are described in further detail below in connection with Fig. 8.

3. Data Transport Management

With reference to Figs. 6, 7 and 8, hardware and/or software components associated with the acquisition 100, broadcasting 102 and reception phases 104 will now be described. These hardware and/or software components comprise various transport components for supporting MOD or live stream content distribution in one or more multicast-enabled networks in the system 10. The transport components can be, but are not limited to, a file transport module, a transport sender, a transport broadcaster, and a transport receiver. The content is preferably characterized as either live content and simulated/scheduled live content, or MOD (i.e., essentially any file). Streaming media such as live content or simulated/scheduled live content are managed and transported similarly, while MOD is handled differently.

Acquisition for plural customers A through X is illustrated in Fig. 6. By way of an example, acquisition for customer A involves an encoder, as indicated at 134, which can employ Real, WMT, MPEG, QT, among other encoding schemes with content from a source 24. The encoder also encodes packets into a format to facilitate broadcasting in accordance with the present invention. A disk 130 stores content from different sources and provides MOD streams, for example, to a disk host 132. The disk host 132 can be proxying the content or hosting it. Live content, teleconferencing, stock and weather data generating systems, and the like, on the other hand, is also encoded. The disk host 132 unicasts the MOD streams to a file transport module 136, whereas the encoder 134 provides the live streams to a transport sender 138 via unicast

or multicast. The encoder can employ either unicast or multicast if QT is used. Conversion from unicast to multicast is not always needed, but multicast-to-multicast conversion can be useful. The file transport module 136 transfers MOD content to a multicast-enabled network. The transport sender 138 pulls stream data from a media encoder 134 or an optional aggregator and sends stream announcements (e.g., using session announcement protocol and session description protocol (SAP/SDP)) and stream data to multicast Internet protocol (IP) addresses and ports received from a transport manager. The transport manager is described below with reference to Fig. 9. When a Real G2 server is used to push a stream, as opposed to a pulling scheme, an aggregator can be used to convert from a push scheme to a pull scheme. The components described in connection with Fig. 6 can be deployed at the encoding center 28 or in a distributed manner at, for example, content provider facilities.

Fig. 7 illustrates an exemplary footprint for one of a plurality of broadcasts. As shown in Fig. 7, the broadcasting phase 102 is implemented using a transport broadcaster 140 and a transport bridge 142. These two modules are preferably implemented as one software program, but different functions, at a master data center 18 or network operations center. The transport broadcaster 140 performs transport path management, whereas the transport bridge 142 provides for peering. The broadcaster 140 and bridge 142 get data from the multicast cloud (e.g., network 108) being guided by the transport manager and forward it to an appropriate transport path. One transport broadcaster 140, for example, can be used to represent one transport path such as satellite uplink or fiber between data centers or even a cross-continental link to a data center in Asia from a data center in North America. The broadcaster 140 and bridge 142 listen to stream announcements from transport senders 138 and enable and disable multicast traffic to another transport path, accordingly. They can also tunnel multicast traffic by using TCP to send stream information and data to another multicast-enabled network. Thus, broadcasters 110 transmit corresponding subsets of the acquisition phase streams that are sent via the multicast cloud 108. In other words, the broadcasters 110 operate as gatekeepers for their respective transport paths, that is, they pass any streams that need to be sent via their corresponding path and prevent passage of other streams. Transmission can also be accomplished using TCP to another

receiver regardless whether the system that the receiver is in is multicast-enabled. Thus, multicast operation can be disabled and the broadcast is still routed and distributed, although not quite as effectively or inexpensively as multicast.

Fig. 8 illustrates the reception phase 104 at one of a plurality of servers or data centers. As stated above, the data centers are preferably deployed in a tiered hierarchy 116, 118 and 120 comprising media serving systems 14, regional data centers 16 and master data centers 18, respectively. The tiers 116, 118 and 120 each comprise a transport receiver 144. Transport receivers can be grouped using, for example, the transport manager. Each transport receiver 144 receives those streams from the broadcasters 110 that are being sent to a group to which the receiver belongs. The transport receiver listens to stream announcements, receives stream data from plural transport senders 138 and feeds the stream data to media servers 146. The transport receiver 144 can also switch streams, as indicated at 154 (e.g., to replace a live stream with a local MOD feed for advertisement insertion purposes). The stream switch 154 can be a plug-in in the Media Server 14 or exist in the server itself to enable switching per end-user 20. The plug-in can interact with an advertisement platform to inject advertisements into streams. The MOD streams are received via the file transport 136 and stored, as indicated via the disk host 148, database 150 and proxy cache/HTTP server 152. The servers 146 and 152 provide content streams to users 20.

The transport components described in connection with Figs. 6, 7 and 8 are advantageous in that they generalize data input schemes from encoders and optional aggregators to data senders, data packets within the system 10, and data feeding from data receivers to media servers, to support essentially any media format. The transport components preferably employ RTP as a packet format and XML-based remote procedure calls (XBM) to communicate between transport components.

The transport manager will now be described with reference to Fig. 9 which illustrates an overview of transport data management. The transport manager is preferably a software module deployed at the encoding facility 28 or other facility designated as a NOC. As shown in Fig. 10, multiple data sources 14 (e.g., database content, programs and applications) provide content as input into the transport manager 170. Information regarding the content from these data sources is also

provided to the transport manager such as identification of input source 14 and output destination (e.g., groups of receivers). Decisions as to where content streams are to be sent and which groups of servers are to receive the streams can be predefined and indicated to the transport manager 170 as a configuration file or XBM function call in
5 real-time. This information can also be entered via a graphical user interface (GUI) 172 or command line utility. In any event, the information is stored in a local database 174. The database 174 also stores information for respective streams relating to defined maximum and minimum IP address and port ranges, bandwidth usage, groups or communities intended to receive the streams, network and stream names, as well as
10 information for user authentication to protect against unauthorized use of streams or other distributed data.

With continued reference to Fig. 9, a customer requests to stream content via the system 10 using, for example, the GUI 172. The request can include the customer's name and account information, the stream name to be published (i.e., distributed) and
15 the IP address and port of the encoder or media server from which the stream can be pulled. Requests and responses are sent via the multicast network (e.g., cloud 108) using separate multicast addresses for each kind of transport component (e.g., a transport sender channel, a broadcaster channel, a transport manager channel and a transport receiver channel), or one multicast address and different ports. IP and port
20 combinations can be used for TCP transmissions. An operator at the NOC 28 can approve the request if sufficient system resources are available such as bandwidth or media server capacity. Automatic approval can be provided by a scheduling system configured to provide immediate responses to attempted broadcasts. The transport manager 170 preferably pulls stream requests periodically. In response to an approved
25 request, the transport manager 170 generates a transport command in response to the request (e.g., an XML-based remote procedure call (XBM)) to the transport sender 138 corresponding to that customer which provides the assigned multicast IP address and port that the transport sender is allowed to use in the system 10.

The transport sender 138 receives the XBM call and responds by announcing the
30 stream that is going to be sent. All of the transport components listen to the announcement. Once the transport sender 138 commences sending the stream into the

assigned multicast IP address and port, the corresponding transport broadcaster 140 filters the stream. The transport receiver 144 joins the multicast IP address and receives the data or stream if the stream is intended for a group to which the receiver 144 belongs. As stated above in connection with Fig. 5, the receiver converts the steam
5 received via the cloud 108 and sends it to the media server available to the users 20. The data is then provided to the media server associated with the receiver. Receivers 144 and broadcasters 140 track announcements that they have honored using link lists.

As stated above, the transport components described with reference to Figs. 5-9 preferably use RPT as a data transport protocol. Accordingly, Windows Media, RealG2
10 and QT packets are wrapped into RTP packets. The acquisition network 22 preferably employs an RTP stack to facilitate processing any data packets, wrapping the data packets with RTP header and sending the data packets. RTSP connection information is generally all that is needed to commence streaming.

RTP is used for transmitting real-time data such as audio and video, and
15 particularly for time-sensitive data such as streaming media, whether transmission is unicast or multicast. RTP employs User Datagram Protocol (UDP), as opposed to Transmission Control Protocol (TCP) that is typically used for non-real-time data such as file transfer and e-mail. Unlike with TCP, software and hardware devices that create and carry UDP packets do not fragment and reassemble them before they have reached
20 their intended destination, which is important in streaming applications. RTP adds header information that is separate from the payload (e.g., content to be distributed) that can be used by the receiver. The header information is merely interpreted as payload by routers that are not configured to use it.

RTSP is an application-level protocol for control over the delivery of data with
25 real-time properties and provides an extensible framework to enable controlled, on-demand delivery of real-time data including live feeds and stored clips. RTSP can control multiple data delivery sessions, provide means for choosing delivery channels such as UDP, multicast UDP and TCP, and provide means for choosing delivery mechanisms based on RTP. HTTP is not suitable for streaming media because it is
30 more of a store-and-forward protocol that is more suitable for web pages and other content that is read repeatedly. Unlike HTTP, RTSP is highly dynamic and provides

persistent interactivity between the user device (hereinafter referred to as a client) and server that is beneficial for time-based media. Further, HTTP does not allow for multiple sessions between a client and server, and travels over only a single port. RTP can encapsulate HTTP data, and can be used to dynamically open multiple RTP sessions to deliver many different streams at the same time.

The system 10 employs transmission control software deployed at the encoding facilities 28, which can operate as a network operations center (NOC), and at broadcasters 110 (e.g., master data centers 120) to determine which streams will be available to which nodes in the distribution system 12 and to enable the distribution system 12 to support one-to-one streaming or one-to-many streaming. The extensible language capabilities of RTSP augment the transmission control software at the edge of the distribution network 12. Since RTSP is a bi-directional protocol, its use enables encoders 134 and receivers 144 to talk to each other, allowing for routing, conditional access (e.g., authentication) and bandwidth control in the distribution network 12. Standard RTSP proxies can be provided between any network components to allow them to communicate with each other. The proxy can therefore manage the RTSP traffic without necessarily understanding the actual content.

For every RTSP stream, there is an RTP stream. Further, RTP sessions support data packing with timestamps and sequence numbers. They can also be used for carrying stereo information, wide screen versions of requested media, different audio tracks, and so on. RTP packets are wrapped in a broadcast protocol. Applications in the receiving phase 104 can use this information to determine when to expect the next packet. Further, system operators can use this information to monitor network 12 and satellite 32 connections to determine the extent of latency, if any.

Encoders and data encapsulators written with RTP as the payload standard are advantageous because off-the-shelf encoders (e.g., MPEG2 encoders) can be introduced without changing the system 10. Further, encoders that output RTP/RTSP can connect to RTP/RTSP transmission servers. In addition, the use of specific encoder and receiver combinations can be eliminated when all of the media players support RTP/RTSP.

4. Testing Component

As stated above, transport management facilitates requests for streaming content by system customers (e.g., content providers). The use of RTP/RTSP between transport components also facilitates communication between servers and users requesting to receive content from the servers.

Media resource requests from a client can occur inside of an RTSP connection or via a simple HTTP request. Responses to media resource requests can be metafiles, but can also occur inside of a binary file or via the protocol being used between the client and server (e.g., server 14, 16 or 18). In each of these instances, responses are similar to a response served by an Internet client-server application to allow sending links to a resource rather than having to send the resource itself. These files and/or response information indicate to the client 22 the location of requested media, i.e., where it should connect to and in what order. In video serving applications, metafiles allow content providers to create playlists of video clips, but metafiles can also be used to help define events and other information such as the author or resource owner. Further, the contents of metafiles can be written as more of a scriptable language to handle conditions and other more dynamic needs. In other words, RTSP can enable encoders 134, receivers 144, and servers or data centers 14, 16 and 18 to communicate with each other to allow for routing, conditional access (e.g., authentication) and bandwidth control in the distribution network 12, in accordance with the present invention. For example, standard RTSP proxies can be provided between any network components to allow them to communicate with each other.

With reference to the illustrative Internet broadcast system 10 for streaming media, the testing component 27 is implemented as a combination of software modules and hardware components in the system 10 and is operable to review properties of the client 20 in real-time using RTP/RTSP communications to determine, for example, the type of client (e.g., which media player is to be used), as well as the bandwidth and capability of the client. Information relating to these properties can be provided in a metafile corresponding to the request. The metafile can also provide information relating to the ISP such as client subscription type and priority to be provided the

requested stream based on contracts between companies using the ISP and content providers. For example, a content provider can pay an ISP for higher quality of service.

The testing component 27 is configurable to analyze requests and place higher priority to the content provider's connections than other requested traffic. The information for authenticating client requests can come from other sources via the internet than the metafile associated with the request, such as from an ISP or other database.

In accordance with the preferred embodiment of the present invention, a single URL is assigned to an Internet resource that is provided in a plurality of formats. In particular, rather than assign a separate URL for each available format of a resource, all formats of a resource are assigned the same URL.

In operation, when a client makes a request for a particular resource using a URL assigned to the resource, the type of resource format which the client can accept is determined by the testing component 27 based on the type of client making the request. The type of format that a particular client or customer can utilize can be known in advance or determined at the time of the request. The type of resource format that the client can accept is then compared to the available resource formats and based on the comparison, the resource is provided/served to the client in a proper format. Accordingly, by determining the type of client making a request for a resource, the system 10 is able to host a resource in multiple formats while still only providing one URL to customers for that resource. The testing component 27 can be implemented in the server 25 in accordance with one embodiment of the present invention.

While a server can host multiple media formats, each format is typically hosted by a different server. In accordance with another aspect of the invention, a client can be re-directed to a selected one of multiple servers, which respectively carry the different formats needed to fulfill a user request using the client information, in order to serve the specific media type needed for the client.

By way of an example, a content provider that wishes to stream media via the system 10 provides the NOC 28 with the URL. The URL is translated at the NOC into publishing points or mount points into selected servers 14, 16 and/or 18 in the distribution system 12. When a client types that URL (e.g., <http://www.ibeam.com/test>), the initial connection is accepted, and the request is

analyzed. If it is determined by the testing component 27 that the request came from a QuickTime player, the testing component 27 engages RTP/RTSP protocol signaling to generate a response that either informs the client to reconnect directly to the QuickTime server for '/test/', or redirects the initial connection to the QuickTime server in a manner similar to a proxy server.

Although the present invention has been described with reference to a preferred embodiment thereof, it will be understood that the invention is not limited to the details thereof. Various modifications and substitutions will occur to those of ordinary skill in the art. All such substitutions are intended to be embraced within the scope of the invention as defined in the appended claims.